

WHAT IS CLAIMED IS:

1. A computer storage system, comprising:

a plurality of disk drives for storing distributed parity groups, each distributed parity group comprising storage blocks, said storage blocks comprising one or more data blocks and a parity block, each of said storage blocks stored on a separate disk drive such that no two storage blocks from a given parity group reside on the same disk drive;

file system metadata to describe a location of each of said storage blocks by specifying a disk identifier and a logical block identifier of a first logical block of each storage block; and

an allocation module configured to allocate space for a new distributed parity group on two or more disk drives in said plurality of disk drives, said allocation module allocating space for each storage block of said new distributed parity group from free space on each of said two or more disk drives.

2. The computer storage system of Claim 1, said allocation module further configured to recognize a new disk drive hot-swapped into said plurality of disk drives during file system operation and to use said new disk drive to store one or more storage blocks.

3. The computer storage system of Claim 1, said allocation module further configured to recognize a new disk drive hot-swapped into said plurality of disk drives during file system operation and to use any free logical blocks on said new disk drive to store one or more storage blocks.

4. The computer storage system of Claim 3, wherein said new disk drive is larger in capacity than any drive in said plurality of disk drives.

5. The computer storage system of Claim 1, wherein a size of a first distributed parity group is larger than a size of a second distributed parity group within a first file.

6. The computer storage system of Claim 1, wherein said new disk drive is provided to a Fibre Channel network.

7. The computer storage system of Claim 1, wherein a file is organized as one or more distributed parity groups of varying size.

8. The computer storage system of Claim 1, wherein an extent of a first distributed parity group of a file is larger than an extent of a second distributed parity group of said file.

5 9. The computer storage system of Claim 1, further comprising a load-balancing module to move one or more distributed parity groups from disk to disk to improve load balancing of said plurality of disks.

10 10. The computer storage system of Claim 1, further comprising a load-balancing module to move one or more distributed parity groups from disk to disk to improve capacity balancing of said plurality of disks.

11 11. The computer storage system of Claim 1, further comprising a load-balancing module to move one or more distributed parity groups from disk to disk to improve load balancing of said plurality of disks when a new disk is added to said plurality of disks.

12 12. The computer storage system of Claim 1, further comprising a load-balancing module to move one or more distributed parity groups from disk to disk to improve capacity balancing of said plurality of disks when a new disk is added to said plurality of disks.

13. A method for managing a storage array, comprising:
recognizing a new storage device;
20 adding said new storage device to a list of previously-available storage devices to produce a list of currently-available storage devices;
moving a selected storage block from a selected parity group to said new storage device;
updating file system metadata to reflect the new location of said selected
25 storage device such that clients accessing a file containing said selected storage block can continue to access data in said selected storage block.

14. The method of Claim 13, wherein said updating comprises specifying a disk identifier and a logical block identifier.

30 15. The method of Claim 13, wherein said new storage device comprises a new disk drive.

